


Streaming Analytics with Spark

Magnoni Luca
IT-CM-MM

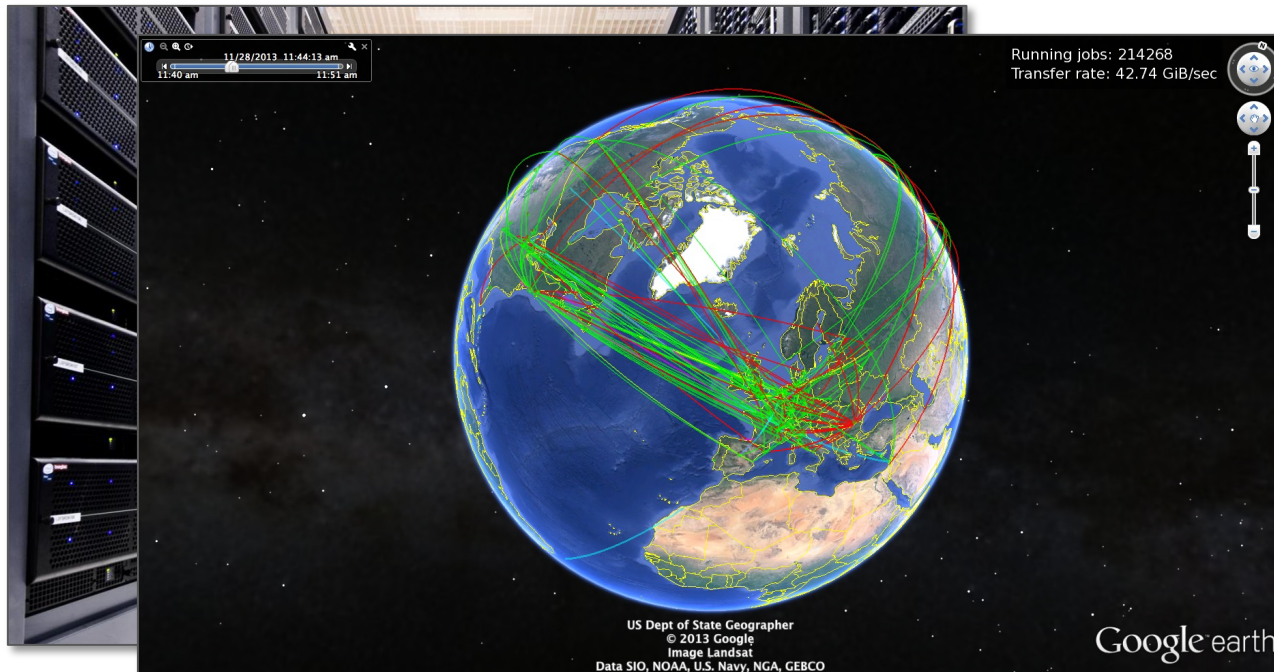
Monitoring @ CERN

- Provide monitoring facilities for CERN Data Centers and WLCG Infrastructure



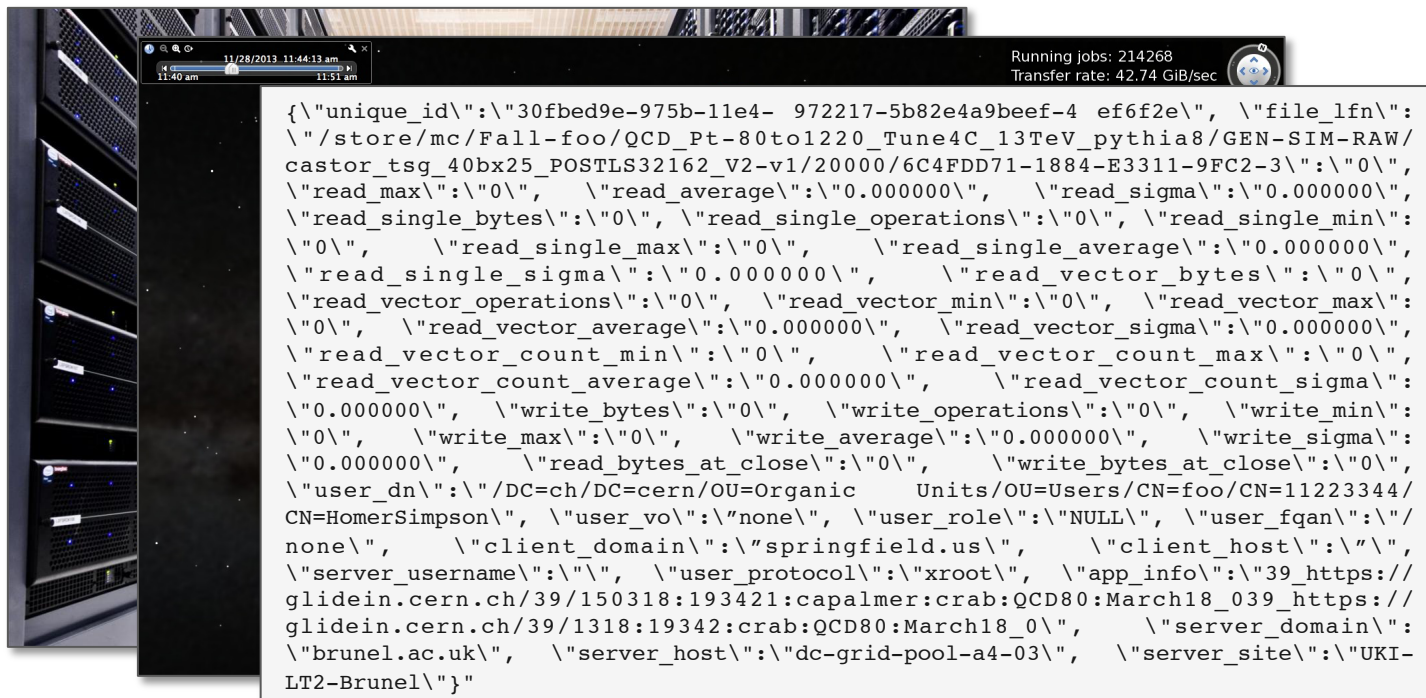
Monitoring @ CERN

- Provide monitoring facilities for CERN Data Centers and WLCG Infrastructure



Monitoring @ CERN

- Provide monitoring facilities for CERN Data Centers and WLCG Infrastructure

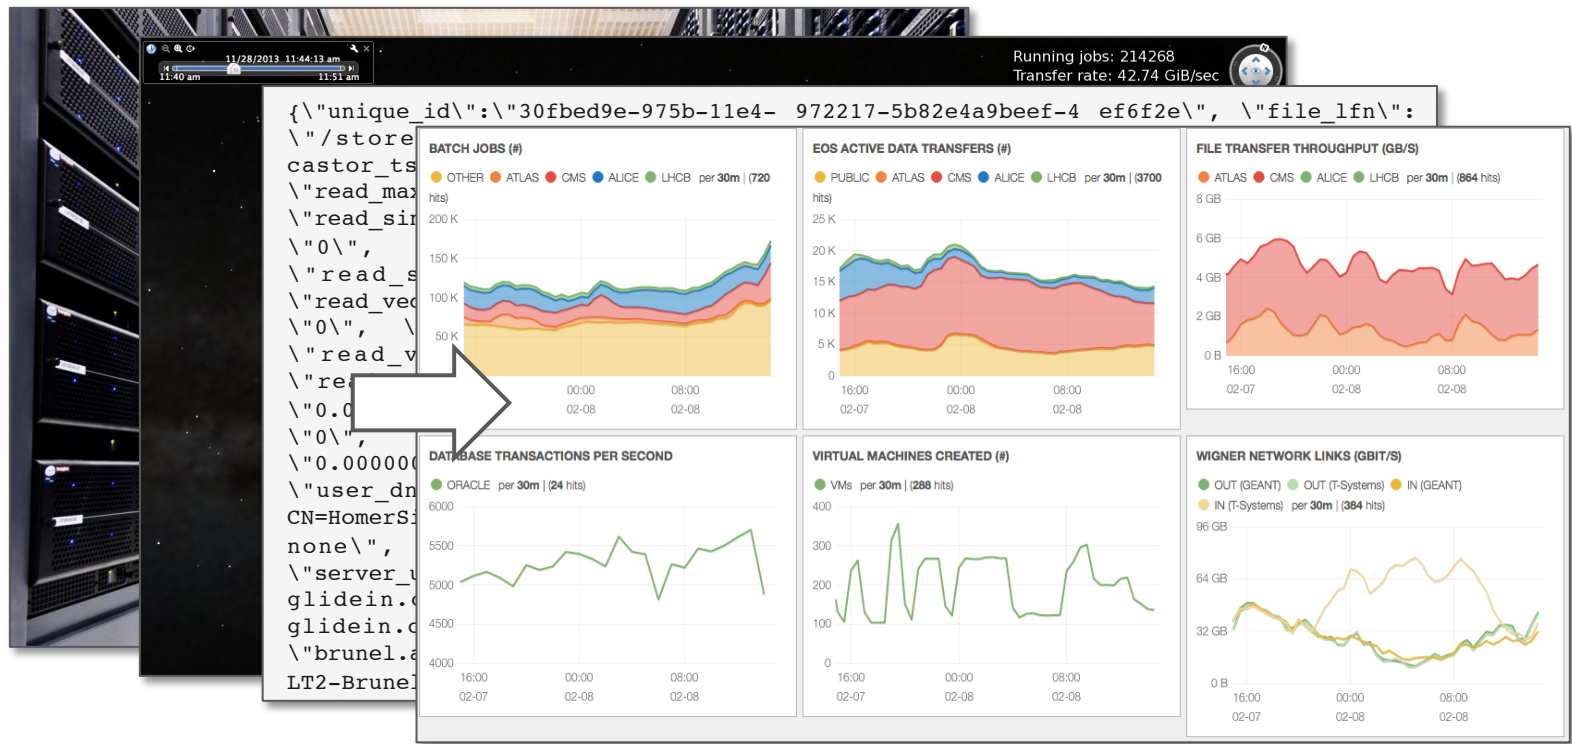


The image shows a server rack on the left and a monitoring window on the right. The window displays a JSON log entry with various performance metrics and file path information. The top of the window shows the date and time as 11/28/2013 11:44:13 am. The bottom right corner of the window displays 'Running jobs: 214268' and 'Transfer rate: 42.74 GiB/sec'.

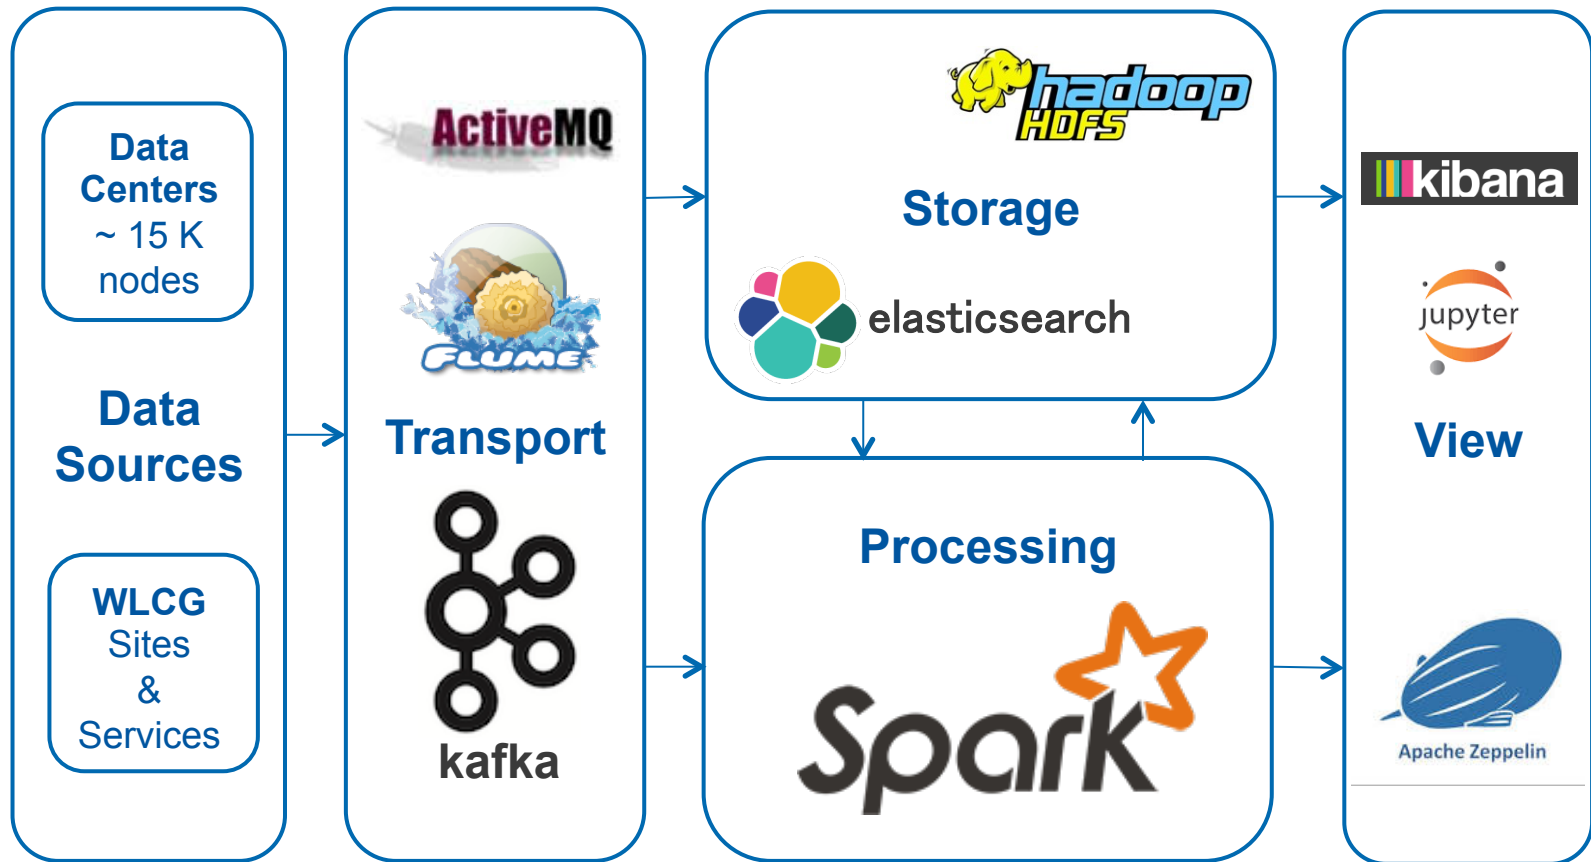
```
{\"unique_id\": \"30fbed9e-975b-11e4- 972217-5b82e4a9beef-4 ef6f2e\", \"file_lfn\": \"/store/mc/Fall-foo/QCD_Pt-80to1220_Tune4C_13TeV_pythia8/GEN-SIM-RAW/castor_tsg_40bx25_POSTLS32162_V2-v1/20000/6C4FDD71-1884-E3311-9FC2-3\": \"0\", \"read_max\": \"0\", \"read_average\": \"0.000000\", \"read_sigma\": \"0.000000\", \"read_single_bytes\": \"0\", \"read_single_operations\": \"0\", \"read_single_min\": \"0\", \"read_single_max\": \"0\", \"read_single_average\": \"0.000000\", \"read_single_sigma\": \"0.000000\", \"read_vector_bytes\": \"0\", \"read_vector_operations\": \"0\", \"read_vector_min\": \"0\", \"read_vector_max\": \"0\", \"read_vector_average\": \"0.000000\", \"read_vector_sigma\": \"0.000000\", \"read_vector_count_min\": \"0\", \"read_vector_count_max\": \"0\", \"read_vector_count_average\": \"0.000000\", \"read_vector_count_sigma\": \"0.000000\", \"write_bytes\": \"0\", \"write_operations\": \"0\", \"write_min\": \"0\", \"write_max\": \"0\", \"write_average\": \"0.000000\", \"write_sigma\": \"0.000000\", \"read_bytes_at_close\": \"0\", \"write_bytes_at_close\": \"0\", \"user_dn\": \"/DC=ch/DC=cern/OU=Organic Units/OU=Users/CN=foo/CN=11223344/CN=HomerSimpson\", \"user_vo\": \"none\", \"user_role\": \"NULL\", \"user_fqan\": \"none\", \"client_domain\": \"springfield.us\", \"client_host\": \"\", \"server_username\": \"\", \"user_protocol\": \"xroot\", \"app_info\": \"39_https://glidein.cern.ch/39/150318:193421:capalmer:crab:QCD80:March18_039_https://glidein.cern.ch/39/1318:19342:crab:QCD80:March18_0\", \"server_domain\": \"brunel.ac.uk\", \"server_host\": \"dc-grid-pool-a4-03\", \"server_site\": \"UKI-LT2-Brunel\"}
```

Monitoring @ CERN

- Provide monitoring facilities for CERN Data Centers and WLCG Infrastructure

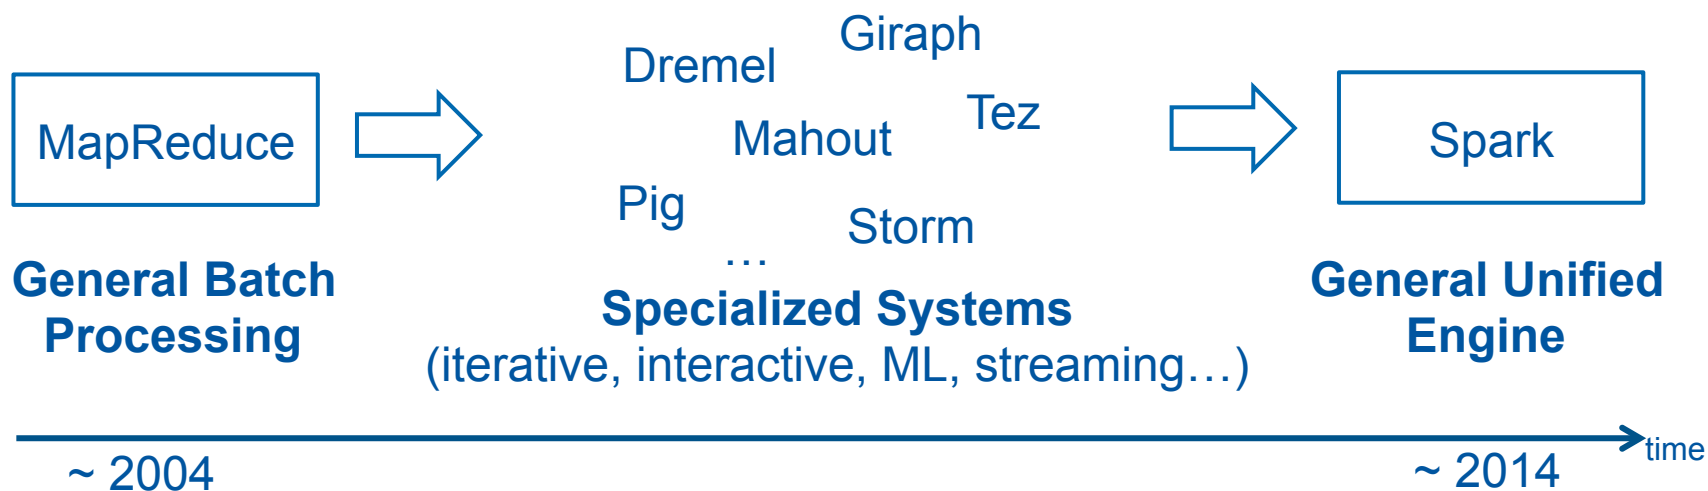


Monitoring Architecture & Tools



Apache Spark

- Distributed large-scale data processing engine
- Most active Apache project in 2014
- Simpler and faster than Hadoop/MapReduce
 - One framework for batch, streaming, SQL, iterative,...



Spark Streaming

- Distributed analysis of data streams
 - Micro-batch computation
- Scalable, High-Throughput, Fault-Tolerant
- Processing latency “*as low as*” few seconds

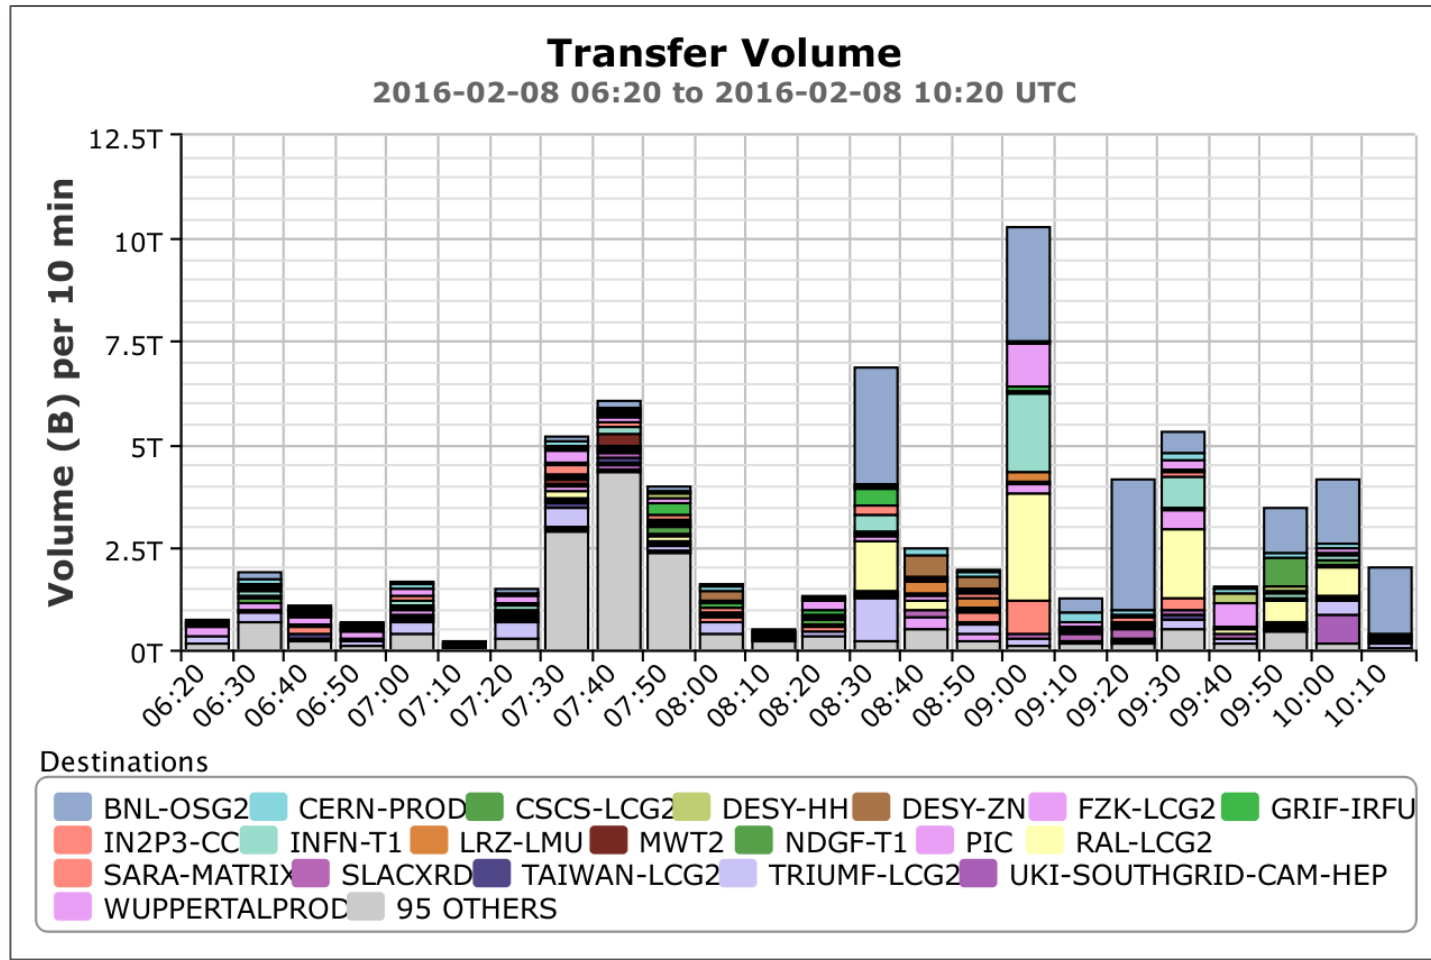


Source: <http://spark.apache.org/docs/latest/streaming-programming-guide.html>

Data Transport for Streaming

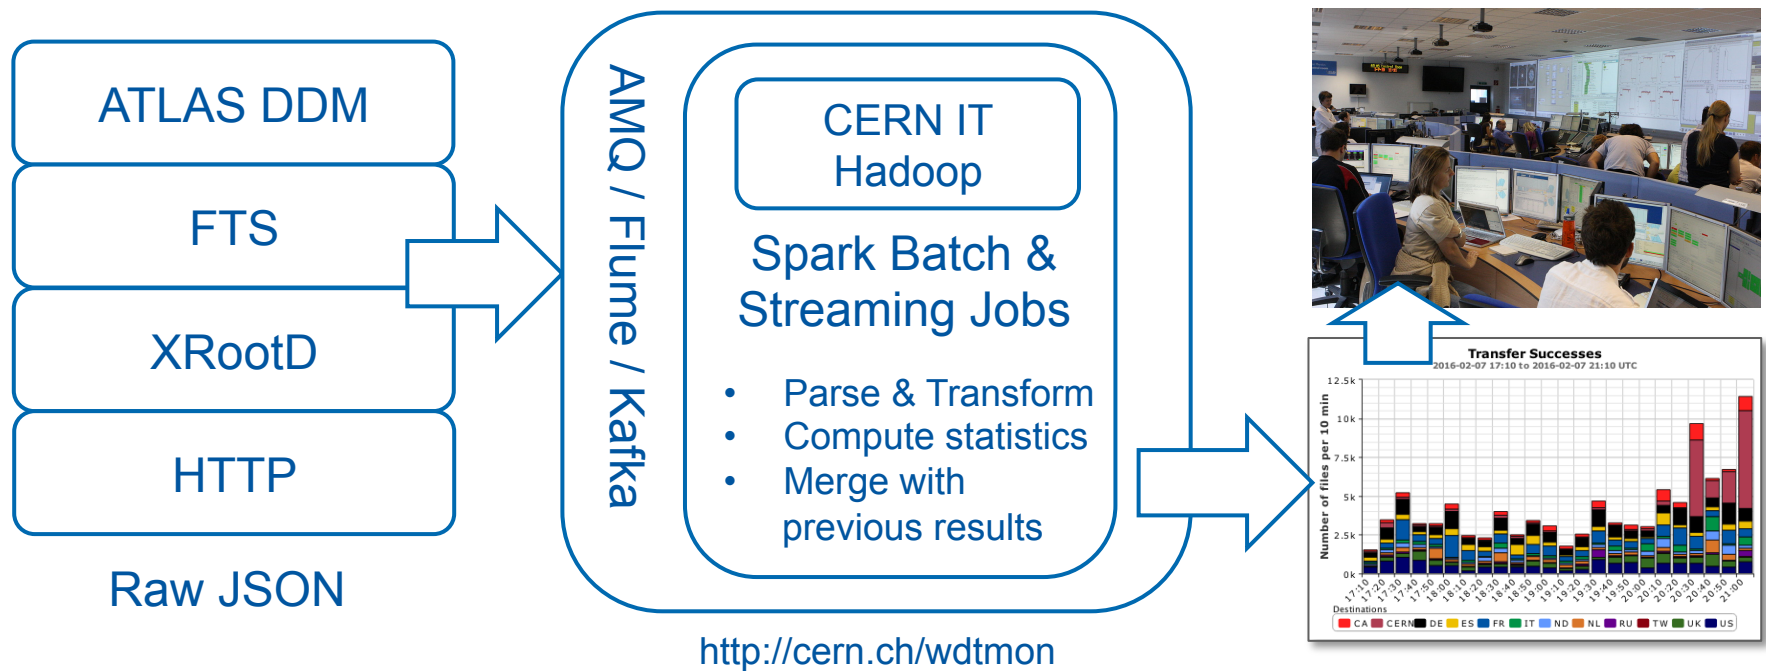
- Streaming relies on transport to serve data as it is produced, at scale & speed
- Tool-chain @ CERN:
 - Collection/Aggregation of metrics and logs
 - **Flume**, one agent on each data centers node
 - Reliable and scalable data consumption
 - **Kafka**, as distributed high-volume publish/subscribe messaging system
 - External producers
 - **ActiveMQ** (AMQ) over STOMP

Spark examples @ CERN: WLCG Data Transfers dashboards



Spark examples @ CERN: WLCG Data Transfers dashboards


- ~ 100 GB/day of monitoring data (data transfers and access logs) to be gathered and processed
- Compute and Aggregate over time the data transfer activities across WLCG sites from raw operation logs




Spark examples @ CERN:

Live analytics

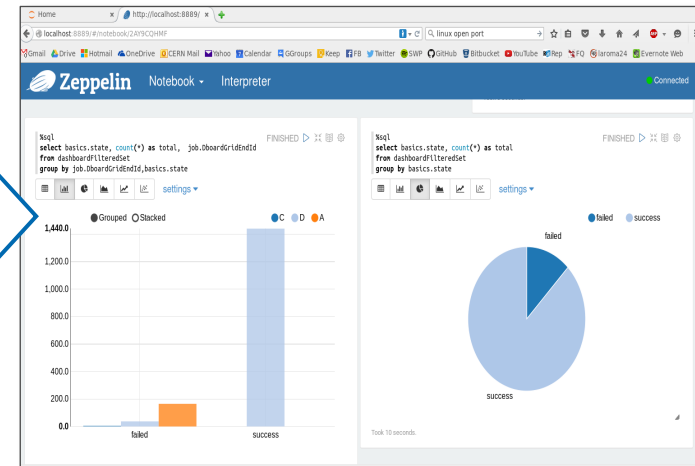

- Explore metrics and logs from data centers as they flow
- Process the raw information to correlate and make sense of the overall behavior (e.g. multi-node alarms)
- Notebooks as a service



Sensors Metrics



<http://foz.cern.ch>



Zeppelin notebook:
Easy and Interactive Analytics

Spark examples @ CERN: *Live analytics*

The screenshot displays the Zeppelin Notebook interface. The top navigation bar includes the Zeppelin logo, 'Notebook', and 'Interpreter' tabs, along with a 'Connected' status indicator. The main workspace is divided into two panels, each showing a SQL query and its visualization.

Left Panel:

```
%sql
select basics.state, count(*) as total, job.DboardGridEndId
from dashboardFilteredSet
group by job.DboardGridEndId, basics.state
```

The visualization is a grouped bar chart. The y-axis represents the count, ranging from 0.0 to 1,440.0. The x-axis shows two categories: 'failed' and 'success'. The 'failed' group has three bars (C, D, A) with counts approximately 10, 50, and 150. The 'success' group has a single bar (C) with a count of 1,440.

Group	State	Count
failed	C	~10
	D	~50
	A	~150
success	C	1,440

Right Panel:

```
%sql
select basics.state, count(*) as total
from dashboardFilteredSet
group by basics.state
```

The visualization is a pie chart showing the distribution of states. The 'success' state (light blue) represents approximately 90% of the total, and the 'failed' state (dark blue) represents approximately 10%.

State	Count
failed	~160
success	1,440

Both panels indicate the query is 'FINISHED' and took 10 seconds to execute.

<http://foz.cern.ch>

Easy and Interactive Analytics

More Spark @ CERN

- IT Analytics Working Group
 - Offline analysis and data exploration
- LHC Experiments activities
 - ATLAS study on data scrutiny and popularity

Thank you!

